

Robots in War: Issues of Risk and Ethics

Patrick LIN^{a,1}, George BEKEY^b, and Keith ABNEY^c

^a*California Polytechnic State University, San Luis Obispo, USA*
United States Naval Academy, Annapolis, USA

^b*California Polytechnic State University, San Luis Obispo, USA*
University of Southern California, Los Angeles, USA

^c*California Polytechnic State University, San Luis Obispo, USA*

Abstract. War robots clearly hold tremendous advantages—from saving the lives of our own soldiers, to safely defusing roadside bombs, to operating in inaccessible and dangerous environments such as mountainside caves and underwater. Without emotions and other liabilities on the battlefield, they could conduct warfare more ethically and effectively than human soldiers who are susceptible to overreactions, anger, vengeance, fatigue, low morale, and so on. But the use of robots, especially autonomous ones, raises a host of ethical and risk issues. This paper offers a survey of such emerging issues in this new but rapidly advancing area of technology.

Keywords. Robot, autonomous, military, war, ethics, risk, responsibility, advantages, benefits, liabilities, harm

Imagine the face of warfare with advanced robotics: Instead of our soldiers returning home in flag-draped caskets to heartbroken parents, autonomous robots – mobile machines that can make decisions, such as to fire upon a target without human intervention—can replace the human soldier in an increasing range of dangerous missions: from tunneling through dark caves in search of terrorists, to securing urban streets rife with sniper fire, to patrolling the skies and waterways where there is little cover from attacks, to clearing roads and seas of improvised explosive devices (IEDs), to surveying damage from biochemical weapons, to guarding borders and buildings, to controlling potentially-hostile crowds, and even as the infantry frontlines.

These robots would be ‘smart’ enough to make decisions that only humans now can; and as conflicts increase in tempo and require much quicker information processing and responses, robots have a distinct advantage over the limited and fallible cognitive capabilities that we *Homo sapiens* have. Not only would robots expand the battlespace over difficult, larger areas of terrain, but they also represent a significant force-multiplier – each effectively doing the work of many human soldiers, while immune to sleep deprivation, fatigue, low morale, perceptual and communication challenges in the ‘fog of war’, and other performance-hindering conditions. But the presumptive case for deploying robots on the battlefield is more than about saving

¹ Contact: Dr. Patrick Lin, Cal Poly, Ethics + Emerging Technologies Group, Philosophy Dept., San Luis Obispo, CA, 93407, USA, palin@calpoly.edu, <http://ethics.calpoly.edu>. This paper is sponsored by the US Department of the Navy, Office of Naval Research, under awards # N00014-07-1-1152 and N00014-08-1-1209. We also thank the College of Liberal Arts and the College of Engineering at Cal Poly for their support, as well as our colleagues for their discussions and contributions to our investigation here, particularly: Colin Allen, Ron Arkin, Peter Asaro, and Wendell Wallach.

human lives or superior efficiency and effectiveness, though saving lives and clearheaded action during frenetic conflicts are significant issues.

Robots, further, would be unaffected by the emotions, adrenaline, and stress that cause soldiers to overreact or deliberately overstep the rules of engagement and commit atrocities, that is to say, war crimes. We would no longer read (as many) news reports about our own soldiers brutalizing enemy combatants or foreign civilians to avenge the deaths of their brothers in arms – unlawful actions that carry a significant political cost. Indeed, robots may act as objective, unblinking observers on the battlefield, reporting any unethical behavior back to command; their mere presence as such would discourage all-too-human atrocities in the first place.

Technology, however, is a double-edge sword with both benefits and risks, critics and advocates; and advanced, autonomous military robotics is no exception, no matter how compelling the case may be to pursue such research. The worries include: where responsibility would fall in cases of unintended or unlawful harm, which could range from the manufacturer to the field commander to even the machine itself; the possibility of serious malfunction and robots gone wild; capturing and hacking of military robots that are then unleashed against us; lowering the threshold for entering conflicts and wars, since fewer military lives would then be at stake; the effect of such robots on squad cohesion, e. g., if robots recorded and reported back the soldier's every action; refusing an otherwise-legitimate order; and other possible harms.

In this paper, we will evaluate these and other concerns, including the driving forces in autonomous military robotics and the need for 'robot ethics.' Though we speak from a US perspective and refer to reports originating from US defense organizations, our discussion may also be applied to advanced military robotics development by any nation.

Introductory Remarks

First, in this investigation, we are not concerned with the question of whether it is even technically possible to make a perfectly ethical robot, i. e., one that makes the 'right' decision in every case or even most cases. Following Arkin, we agree that an ethically-infallible machine ought not to be the goal now (if it is even possible); rather, our goal should be more practical and immediate: to design a machine that performs better than humans do on the battlefield, particularly with respect to reducing unlawful behavior or war crimes (Arkin, 2007). Considering the number of incidences of unlawful behavior – and by 'unlawful' we mean a violation of the various laws of war (LOW; also referred to elsewhere as Laws of Armed Conflict or LOAC) or rules of engagement (ROE), such as established by the Geneva and Hague Conventions – this appears to be a low standard to satisfy, though a profoundly important hurdle to clear. To that end, scientists and engineers need not first solve the daunting task of creating a truly 'ethical' robot, at least in the foreseeable future; rather, it seems that they only need to program a robot to act in compliance with the LOW and ROE (though this may not be as straightforward and simple as it first appears) or act ethically in limited situations.

Second, we should note that the purpose of this paper is not to limit research with autonomous military robotics, but rather to help responsibly guide it. That there should be two faces to technology – benefits and risk – is not surprising, as history shows, and

is not by itself an argument against that technology.² But ignoring those risks, or at least only reactively addressing them and waiting for public reaction, seems to be unwise, given that it can lead (and, in the case of biotech foods, has led) to a backlash that stalls forward progress.

That said, it is surprising to note that one of the most comprehensive and recent reports on military robotics, *Unmanned Systems Roadmap 2007-2032*, does not mention the word 'ethics' once nor risks raised by robotics, with the exception of one sentence that merely acknowledges that "privacy issues [have been] raised in some quarters" without even discussing said issues (US Department of Defense, 2007, p. 48). While this omission may be understandable from a public relations standpoint, again it seems short-sighted given lessons in technology ethics, especially from our recent past. Our investigation, then, is designed to address that gap, proactively and objectively engaging policymakers and the public to head off a potential backlash that serves no one's interests.

Third, while this paper focuses on issues related to advanced and autonomous military robotics, the discussion may also apply well and overlap with issues related to advanced and autonomous military systems, i. e., computer networks. Further, we are focusing on battlefield applications, as opposed to robotics in manufacturing or medicine even if they are supported by military programs, for several reasons as follow. The most contentious military robots will be the weaponized ones: "Weaponized unmanned systems is a highly controversial issue that will require a patient 'crawl-walk-run' approach as each application's reliability and performance is proved" (US Department of Defense, 2007, p. 54). Their deployment is inherently about human life and death, both intended and unintended, so they immediately raise serious concerns related to ethics (e. g., does just-war theory or the LOW/ROE allow for deployment of autonomous fighting systems in the first place?) as well as risk (e. g., malfunctions and emergent, unexpected behavior) that demand greater attention than other robotics applications.

Finally, though a relatively small number of military personnel is ever exposed on the battlefield, loss of life and property during armed conflict has non-trivial political costs, never mind environmental and economic costs, especially if 'collateral' or unintended damage is inflicted and even more so if it results from abusive, unlawful behavior by our own soldiers. How a nation prosecutes a war or conflict receives particular scrutiny from the media and public, whose opinions influence military and foreign policy even if those opinions are disproportionately drawn from events on the battlefield, rather than on the many more developments outside the military theater. Therefore, though autonomous battlefield or weaponized robots may be years away and account for only one segment of the entire military robotics population, there is much practical value in sorting through their associative issues sooner rather than later.

² Biotechnology, for instance, promises to reduce world hunger by promoting greater and more nutritious agricultural and livestock yield; yet continuing concerns about the possible dissemination of bio-engineered seeds (or 'Frankenfoods') into the wild, displacing native plants and crops, have prompted the industry to move more cautiously (e. g., Thompson, 2007). Even Internet technologies, as valuable as they have been in connecting us to information, social networks, etc., and in making new ways of life possible, reveal a darker world of online scams, privacy violations, piracy, viruses, and other ills; yet no one suggests that we should do away with cyberspace (e. g., Weckert, 2007).

1. Market Forces and Considerations

Several industry trends and recent developments – including high-profile failures of semi-autonomous systems, as perhaps a harbinger of challenges with more advanced systems – highlight the need for a technology risk assessment, as well as a broader study of other ethical and social issues related to the field. In the following, we will briefly discuss seven primary market forces that are driving the development of military robotics as well as the need for a guiding ethics; these roughly map to what have been called ‘push’ (technology) and ‘pull’ (social and cultural) factors (US Department of Defense, 2007, p. 44).

1.1. *Compelling Military Utility*

Defense organizations worldwide are attracted to the use of robots for a range of benefits, some of which we have mentioned above. A primary reason is to replace us less-durable humans in “dull, dirty, and dangerous” jobs (US Department of Defense, 2007, p. 19). This includes: extended reconnaissance missions, which stretch the limits of human endurance to its breaking point; environmental sampling after a nuclear or biochemical attack, which had previously led to deaths and long-term effects on the surveying teams; and neutralizing IEDs, which have caused over 40% of US casualties in Iraq since 2003 (Iraq Coalition Casualty Count, 2008). While official statistics are difficult to locate, news organizations report that the US has deployed over 5,000 robots in Iraq and Afghanistan, which have neutralized 10,000 IEDs by 2007 (CBS, 2007).

Also mentioned above, military robots may be more discriminating, efficient, and effective. Their dispassionate and detached approach to their work could significantly reduce the instances of unethical behavior in wartime – abuses that negatively color the prosecution of a conflict, no matter how just the initial reasons to enter the conflict are, and carry a high political cost.

1.2. *US Congressional Deadlines*

Clearly, there is a tremendous advantage to employing robots on the battlefield, and the US government recognizes this. Two key Congressional mandates are driving the use of military robotics in the US: by 2010, one-third of all operational deep-strike aircraft must be unmanned, and by 2015, one-third of all ground combat vehicles must be unmanned (National Defense Authorization Act, 2000). Most, if not all, of the robotics in use and under development are semi-autonomous at best; and though the technology to (responsibly) create fully autonomous robots is near but not quite in hand, we would expect the US Department of Defense to adopt the same, sensible ‘crawl-walk-run’ approach as with weaponized systems, given the serious inherent risks.

Nonetheless, these deadlines apply increasing pressure to develop and deploy robotics, including autonomous vehicles; yet a ‘rush to market’ increases the risk for inadequate design or programming. Worse, without a sustained and significant effort to build in ethical controls in autonomous systems, or even to discuss the relevant areas of ethics and risk, there is little hope that the early generations of such systems and robots will be adequate, making mistakes that may cost human lives.

1.3. Continuing Unethical Battlefield Conduct

Beyond popular news reports and images of purportedly unethical behavior by human soldiers, the US Army Surgeon General's Office had surveyed US troops in Iraq on issues in battlefield ethics and discovered worrisome results. From its summary of findings, among other statistics: "Less than half of Soldiers and Marines believed that non-combatants should be treated with respect and dignity and well over a third believed that torture should be allowed to save the life of a fellow team member. About 10% of Soldiers and Marines reported mistreating an Iraqi non-combatant when it wasn't necessary...Less than half of Soldiers and Marines would report a team member for unethical behavior...Although reporting ethical training, nearly a third of Soldiers and Marines reported encountering ethical situations in Iraq in which they didn't know how to respond" (US Army Surgeon General's Office, 2006). The most recent survey by the same organization reported similar results (US Army Surgeon General's Office, 2008).

Wartime atrocities have occurred since the beginning of human history, so we are not operating under the illusion that they can be eliminated altogether (nor that armed conflicts can be eliminated either, at least in the foreseeable future). However, to the extent that military robots can considerably reduce unethical conduct on the battlefield – greatly reducing human and political costs – there is a compelling reason to pursue their development as well as to study their capacity to act ethically.

1.4. Military Robotics Failures

More than theoretical problems, military robotics have already failed on the battlefield, creating concerns with their deployment (and perhaps even more concern for more advanced, complicated systems) that ought to be addressed before speculation, incomplete information, and hype fill the gap in public dialogue.

In April 2008, several TALON SWORDS units – mobile robots armed with machine guns – in Iraq were reported to be grounded for reasons not fully disclosed, though early reports claim the robots, without being commanded to, trained their guns on 'friendly' soldiers (e. g., Page, 2008); and later reports denied this account but admitted there had been malfunctions during the development and testing phase prior to deployment (e. g., Sofge, 2008). The full story does not appear to have yet emerged, but either way, the incident underscores the public's anxiety – and the military's sensitivity – with the use of robotics on the battlefield (also see 'Public perceptions' below).

Further, it is not implausible to suggest that these robots may fail, because it has already happened elsewhere: in October 2007, a semi-autonomous robotic cannon deployed by the South African army malfunctioned, killing nine 'friendly' soldiers and wounding 14 others (e. g., Shachtman, 2007). Communication failures and errors have been blamed for several unmanned aerial vehicle (UAV) crashes, from those owned by the Sri Lanka Air Force to the US Border Patrol (e. g., BBC, 2005; National Transportation Safety Board, 2007). Computer-related technology in general is especially susceptible to malfunctions and 'bugs' given their complexity and even after many generations of a product cycle; thus, it is reasonable to expect similar challenges with robotics.

1.5. Related Civilian Systems Failures

On a similar technology path as autonomous robots, civilian computer systems have failed and raised worries that can carry over to military applications. For instance, such civilian systems have been blamed for massive power outages: in early 2008, the US state of Florida suffered through massive blackouts across the entire state, as utility computer systems automatically shut off and rerouted power after just a small fire caused by a failed switch at one electrical substation (e. g., Padgett, 2008); and in the summer 2003, a single fallen tree had triggered a tsunami of cascading computer-initiated blackouts that affected tens of millions of customers for days and weeks across the eastern US and Canada, leaving practically no time for human intervention to fix what should have been a simple problem of stopping the disastrous chain reaction (e. g., US Department of Energy, 2004). Thus, it is a concern that we also may not be able to halt some (potentially-fatal) chain of events caused by autonomous military systems that process information and can act at speeds incomprehensible to us, e. g., with high-speed unmanned aerial vehicles.

Further, civilian robotics are becoming more pervasive. Never mind seemingly-harmless entertainment robots, some major cities (e. g., Atlanta, London, Paris, Copenhagen) already boast driverless transportation systems, again creating potential worries and ethical dilemmas (e. g., bringing to life the famous thought-experiment in philosophy: should a fast-moving train divert itself to another track in order to kill only one innocent person, or continue forward to kill the five on its current path?). So there can be lessons for military robotics that can be transferred from civilian robotics and automated decision-making, and vice versa. Also, as robots become more pervasive in the public marketplace – they are already abundant in manufacturing and other industries – the broader public will become more aware of risk and ethical issues associated with such innovations, concerns that inevitably will carry over to the military's use.

1.6. Complexity and Unpredictability

Perhaps robot ethics has not received the attention it needs, given a common misconception that robots will do only what we have programmed them to do. Unfortunately, such a belief is a sorely outdated, harking back to a time when computers were simpler and their programs could be written and understood by a single person. Now, programs with millions of lines of code are written by teams of programmers, none of whom knows the entire program; hence, no individual can predict the effect of a given command with absolute certainty, since portions of large programs may interact in unexpected, untested ways. (And even straightforward, simple rules such as Asimov's Laws of Robotics (Asimov, 1950) can create unexpected dilemmas.) Furthermore, increasing complexity may lead to emergent behaviors, i. e., behaviors not programmed but arising out of sheer complexity (e. g., Kurzweil, 1999, 2005).

Related major research efforts also are being devoted to enabling robots to learn from experience, raising the question of whether we predict with reasonable certainty what the robot will learn. The answer seems to be negative, since if we could predict that, we would simply program the robot in the first place, instead of requiring learning. Learning may enable the robot to respond to novel situations, given the impracticality and impossibility of predicting all eventualities on the designer's part. Thus,

unpredictability in the behavior of complex robots is a major source of worry, especially if robots are to operate in unstructured environments, rather than the carefully-structured domain of a factory.

1.7. Public Perceptions

From Asimov's science fiction novels to Hollywood movies such as "Wall-E", "Iron Man", "Transformers", "Blade Runner", "Star Wars", "Terminator", "Robocop", "2001: A Space Odyssey", and "I, Robot" (to name only a few, from the iconic to recently released), robots have captured the global public's imagination for decades now. But in nearly every one of those works, the use of robots in society is in tension with ethics and even the survival of humankind. The public, then, is already sensitive to the risks posed by robots – whether or not those concerns are actually justified or plausible – to a degree unprecedented in science and technology. Now, technical advances in robotics is catching up to literary and theatrical accounts, so the seeds of worry that have long been planted in the public consciousness will grow into close scrutiny of the robotics industry with respect to those ethical issues, e. g., the book "Love and Sex with Robots" (Levy, 2007) that reasonably anticipates human-robot relationships.

2. The Issues

From the preceding, it should be clear that there are myriad issues in risk and ethics related to autonomous military robotics. Here, we loosely organize these challenges in thematic sub-groups: legal, just war, technical, robot-human, societal, and other and future challenges. This is not meant to be an exhaustive list, as other issues certainly will emerge as the technology develops and field use broadens.³ Nor is this meant to be a substantive discussion – particularly since no easy answers are apparent – but rather it is a broad survey of the issues; each one will require a more detailed investigation.⁴

2.1. Legal Challenges

2.1.1. Unclear Responsibility

To whom would we assign blame – and punishment – for improper conduct and unauthorized harms caused by an autonomous robot (whether by error or intentional): the designers, robot manufacturer, procurement officer, robot controller/supervisor, field commander, a nation's president or prime minister...or the robot itself (Asaro, 2007; Sparrow, 2007; Sharkey, 2008)?

³ As an example of an unexpected policy change, when German forces during World War II recognized the impracticality of using naval submarines to rescue crews of sinking enemy ships—given limited space inside the submarine as well as exposure to radar and attacks when they surface—they issued the Laconia Order in 1942, based on military necessity, that released submarines from a long-standing moral obligation for sea vessels to rescue survivors; other nations soon followed suit to effectively eliminate the military convention altogether (Walzer, 1977, pp. 147-151).

⁴ For a discussion of military robotics technology, please see *Unmanned Systems Roadmap 2007-2032* (US Department of Defense, 2007); and for a discussion of autonomous robotics in general, please see *Autonomous Robots: From Biological Inspiration to Implementation and Control* (Bekey, 2005).

In a military system, it may be possible to simply stipulate a chain of responsibility, e. g., the commanding officer is ultimately responsible. But this may oversimplify matters, e. g., inadequate testing allowed a design problem to slip by and caused the improper robotic behavior, in which case perhaps a procurement officer or the manufacturer ought to be responsible. The situation becomes much more complex and interesting with robots that have greater degrees of autonomy, which may make it appropriate to treat them as quasi-persons, if not full moral agents some point in the future. We note that Kurzweil forecasts that, by the year 2029, “[m]achines will claim to be conscious and these claims will be largely accepted” (Kurzweil, 1999).

2.1.2. Refusing an Order

A conflict may arise in the following situation, among others: A commander orders a robot to attack a house that is known to harbor insurgents, but the robot – being equipped with sensors to ‘see’ through walls – detects many children inside and, given its programmed instruction (based on the ROE) to minimize civilian casualties, refuses the order. How ought the situation proceed: should we defer to the robot who may have better situational awareness, or the officer who (as far as she or he knows) issues a legitimate command? This dilemma also relates back to the question of responsibility: if the robot refuses an order, then who would be responsible for the events that ensue? Following legitimate orders is clearly an essential tenet for military organizations to function, but if we permit robots to refuse an order, this may expand the circumstances in which human soldiers may refuse an order as well (for better or worse).

2.1.3. Consent by Soldiers to Risks

We already mentioned the semi-autonomous robotic cannon deployed by the South African army malfunctioned, killing nine ‘friendly’ soldiers and wounding 14 others. It would be naive to think such accidents will not happen again. In these cases, should soldiers be informed that an unusual or new risk exists, e. g., when they are handling or working with other dangerous items, such as explosives or even anthrax? Does consent to risk matter anyway, if soldiers generally lack the right to refuse a work order?

2.2. Just-War Challenges

2.2.1. Attack Decisions

It may be important for the above issue of responsibility to decide who, or what, makes the decision for a robot to strike. Some situations may develop so quickly and require such rapid information processing that we would want to entrust our robots and systems to make critical decisions. But the LOW and ROE demand there to be ‘eyes on target’, either in-person or electronically and presumably in real time. (This is another reason why there is a general ban on landmines: without eyes on target, we do not know who is harmed by the ordnance and therefore have not fulfilled our responsibility to discriminate combatants from non-combatants.) If human soldiers must monitor the actions of each robot as they occur, this may limit the effectiveness for which the robot was designed in the first place: robots may be deployed precisely because they can act more quickly, and with better information, than humans can.

However, some military robots – such as the US Navy’s Phalanx CIWS – seem to already and completely operate autonomously, i. e., they make attack decisions without

human eyes on target or approval. This raises the question of how strictly we should take the ‘eyes on target’ requirement. One plausible argument for stretching that requirement is that the Phalanx CIWS operates as a last line of defense against imminent threats, e. g., incoming missiles in the dark of the night, so the benefits more clearly outweigh the risks in such a case. Another argument perhaps would be that ‘eyes on target’ need not be *human* eyes, whether directly or monitoring the images captured by a remote camera; that is, a human does not necessarily need to directly confirm a target or authorize a strike. A robot’s target-identification module – assuming it has been sufficiently tested for accuracy – programmed by engineers is arguably a proxy for human eyes. At least this gives the system some reasonable ability to discriminate among targets, in contrast to a landmine, for instance. A requirement for 100% accuracy in target identification may be overly burdensome, since that is not a bar we can meet with human soldiers.

2.2.2. Lower Barriers for War

Does the use of advanced weaponry such as autonomous robotics make it easier for one nation to engage in war or adopt aggressive foreign (and domestic) policies that might provoke other nations? If so, is this a violation of *jus ad bellum*—the conditions under which a state may morally engage in warfare (Asaro, 2008; Kahn, 2002)? It may be true that new strategies, tactics, and technologies make armed conflict an easier path to choose for a nation, if they reduce risks to our side. Yet while it seems obvious that we should want to reduce casualties from our respective nations, there is something sensible about the need for some terrible cost to war as a deterrent against entering war in the first place. This is the basis for just-war theory, that war ought to be the very last resort given its horrific costs (Walzer, 1977).

But the considered objection – that advanced robotics immorally lowers barriers for war – hides a logical implication that we should not do anything that makes armed conflict more palatable: we should not attempt to reduce friendly casualties, or improve battlefield medicine, or conduct any more research that would make victory more likely and quicker. Taken to the extreme, the objection seems to imply that we should raise barriers to war, to make fighting as brutal as possible (e. g., using primitive weapons without armor) so that we would never engage in it unless it were truly the last resort. Such a position appears counterintuitive at best and dangerously foolish at worst, particularly if we expect that other nations would not readily adopt a policy of relinquishment, which would put the nation that forgoes advanced robotics at a competitive disadvantage.

2.2.3. Imprecision in LOW and ROE

Asimov’s Laws appear to be as simple as programmable rules can be for autonomous robots, yet they yielded surprising, unintended implications in his stories (e. g., Asimov, 1950). Likewise, we may understand each rule of engagement and believe them to be sensible, but are they truly consistent with one another and sufficiently clear—which appears to be a requirement in order for them to be programmable? Much more complex than Asimov’s Laws, the LOW and ROE leave much room for contradictory or vague imperatives, which may result in undesired and unexpected behavior in robots.

For instance, the ROE to minimize collateral damage is vague: is the rule that we should not attack a position if civilian deaths are expected to be greater than – or even half of – combatant deaths? Are we permitted to kill one (high-ranking) combatant,

even if it involves the death of five civilians – or \$10M in unnecessary damage? A robot may need specific numbers to know exactly where this line is drawn, in order to comply with the ROE. Unfortunately, this is not an area that has been precisely quantified nor easily lends itself for such a determination.

2.3. Technical Challenges

2.3.1. Discriminating among Targets

Some experts contend that it is simply too difficult to design a machine that can distinguish between a combatant and a non-combatant, particularly as insurgents pose as civilians, as required for the LOW and ROE (e. g., Sharkey, 2008; Sparrow, 2007; Canning et al., 2004). Further, robots would need to discriminate between active combatants and wounded ones who are unable to fight or have surrendered. Admittedly, this is a complex technical task, but we need to be clear on how accurate this discrimination needs to be. That is, discrimination among targets is also a difficult, error-prone task for human soldiers, so ought we hold machines to a higher standard than we have yet to achieve ourselves, at least in the near term?

Consider the following: A robot enters a building known to harbor terrorists, but at the same time an innocent girl is running toward the robot (unintentionally) in chasing after a ball that happens to be rolling in the direction of the robot. Would the robot know to stand down and not attack the child? If the robot were to attack, of course that would cause outrage from opposing forces and even our own media and public; but this scenario could likely be the same as with a human soldier, adrenaline running high, who may misidentify the charging target as well. It seems that in such a situation, a robot may be less likely to attack the child, since the robot is not prone to overreact from the influence of emotions and fear, which afflict human soldiers. But in any event, if a robot would likely not perform worse than a human soldier, perhaps this is good enough for the moment until the technical ability to discriminate among targets improves. Some critics, however, may still insist on perfect discrimination or at least far better than humans are capable of, though it is unclear why we should hold robots to such a high standard before such a technology exists (unless their point is to not use robots at all until we have perfected them, which is also a contentious position).

2.3.2. First-Generation Problem

We previously mentioned that it would be naive to believe that another accident with military robots will not happen again. As with any other technologies, errors or bugs will inevitably exist, which can be corrected in the next generation of the technology. With Internet technologies, for instance, first-generation mistakes are not too serious and can be fixed with software patches or updates. But with military robotics, the stakes are much higher, since human lives may be lost as a result of programming or other errors. So it seems that the prudent or morally correct course of action is to rigorously test the robot before deploying it.

However, testing already occurs with today's robots, yet it is still difficult if not impossible to certify any given robot as error-free, given that (a) testing environments may be substantially different than more complex, unstructured, and dynamic battlefield conditions in which we cannot anticipate all possible contingencies; and (b) the computer program used in the robot's on-board computer (its 'brain') may consist of millions of lines of code.

Beta-testing of a program (testing prior to the official product launch, whether related to robotics, business applications, etc.) is conducted today, yet new errors are routinely found in software by actual users even after its official product launch. It is simply not possible to run a complex piece of software through all possible uses in a testing phase; surprises may occur during its actual use. Likewise, it is not reasonable to expect that testing of robots will catch any and all flaws; the robots may behave in unexpected and unintended ways during actual field use. Again, the stakes are high with deploying robots, since any error could be fatal. This makes the first-generation problem, as well as ongoing safety and dependability, an especially sensitive issue (e.g., Van der Loos, 2007).

2.3.3. Robots Running Amok

As depicted in science-fiction novels and movies, some imagine the possibility that robots might break free from their human programming through methods as: their own learning, or creating other robots without such constraints (self-replicating and self-revising), or malfunction, or programming error, or even intentional hacking (e.g., Joy, 2000). In these scenarios, because robots are built to be durable and even with attack capabilities, they would be extremely difficult to defeat – which is the point of using robots as force multipliers. Some of these scenarios are more likely than others: we wouldn't see the ability of robots to fully manufacture other robots or to radically evolve their intelligence and escape any programmed morality for quite some time. But other scenarios, such as hacking, seem to be near-term possibilities, especially if robots are not given strong self-defense capabilities (see below).

That robots might run amok is an enhanced version of the worry that enemies might use our own creations against us, but it also introduces a new element in that previous weapon systems still need a human operator which is a point of vulnerability, i.e., a 'soft underbelly' of the system. Autonomous robots would be designed to operate without human control. What precautions can be taken to prevent one from being captured and reverse-engineered or reprogrammed to attack our own forces? If we design a 'kill switch' that can automatically shut off a robot, this may present a key vulnerability that can be exploited by the enemy.

2.3.4. Unauthorized Overrides

This concern is similar to that with nuclear weapons: that a rogue officer may be enough to take control of these terrible weapons and unleash them without authorization or otherwise override their programming to commit some unlawful action. This is a persistent worry with any new, devastating technology and is a multi-faceted challenge: it is a human problem (to develop ethical, competent officers), an organizational problem (to provide procedural safeguards), and technical problem (to provide systemic safeguards). So there does not yet appear to be anything unique about this worry that should hinder the development or deployment of advanced robotics, to the extent that the concern does not impact the development of other technologies. But it nevertheless is a concern that needs to be considered in the design and deployment phases.

2.3.5. Competing Ethical Frameworks

If we seek to build an ethical framework for action in robots, it is not clear which ethical theory to use as our model (e. g., Anderson and Anderson, 2007). Any sophisticated theory – including deontological ethics, consequentialism, virtue ethics, et al. – seems to be vulnerable to inconsistencies and competing directives (especially if a three- or four-rule system as simple as Asimov's cannot work perfectly); thus a hybrid theory may be required (Wallach and Allen, 2008). This concern is related to the first technical challenge described here, that it is too difficult to embed these behavioral rules or programming into a machine. But we should recall our stated mission here: our goal ought not be to create a perfectly ethical robot, only one that acts more ethically than humans—and sadly this may be a low hurdle to clear.

2.3.6. Coordinated Attacks

Generally, it is better to have more data than less when making decisions, particularly one as weighty as a military strike decision. Robots can be designed to easily network with other robots and systems; but this may complicate matters for robot engineers as well as commanders. We may need to establish a chain of command within robots when they operate as a team, as well as ensure coordination of their actions. The risk here is that as complexity of any system increases, the more opportunities exist for errors to be introduced, and again mistakes by military robots may be fatal.

2.4. Human-Robot Challenges

2.4.1. Effect on Squad Cohesion

As a 'band of brothers', there understandably needs to be strong trust and support among soldiers, just as there is among police officers, firefighters, and so on. But sometimes this sense of camaraderie can be overdeveloped to the extent that one team member becomes complicit in or deliberately assists in covering up an illegal or inappropriate action of another team member. We have discussed the benefits of military robots with respect to behavior that is more ethical than currently exhibited by human soldiers. But robots will also likely be equipped with video cameras and other such sensors to record and report actions on the battlefield. This could negatively impact the cohesion among team or squad members by eroding trust with the robot as well as among fellow soldiers who then may or may not support each other as much anymore, knowing that they are being watched. Of course, soldiers and other professionals should not be giving each other unlawful 'support' anyway; but there may be situations in which a soldier is unclear about or unaware of motivations, orders, or other relevant details and err on the side of caution, i. e., not providing support even when it is justified and needed.

2.4.2. Self-Defense

Asimov's Laws permitted robots to defend themselves where that action did not conflict with higher duties, i. e., harm humans (or humanity) or conflict with a human-issued order. But Arkin suggests that military robots can be more conservative in their actions, i. e., hold their fire, because they do not have a natural instinct of self-preservation and may be programmed without such (Arkin, 2007). But how practical is it, at least economically speaking, to not give robots – which may range from \$100,000

to millions of dollars in cost – the ability to defend itself? If a person, say, a civilian, threatens to destroy a robot, shouldn't it have the ability to protect itself, our very expensive taxpayer-funded investment?

Further, self-defense capabilities may be important for the robot to elude capture and hacking, as previously discussed. Robots may be easily trapped and recovered fully intact, unlike tanks and aircraft, for instance, which usually sustain much if not total damage in order to capture it. These considerations are in tension with using robots for a more ethical prosecution of war, since a predilection to hold their fire would be a major safeguard against accidental fatalities, e. g., mistakenly opening fire on non-combatants; therefore, a tradeoff or compromise among these goals – to have a more ethical robot and to protect the robot from damage and capture – may be needed.

2.4.3. Winning Hearts and Minds

Just-war theory, specifically *jus post bellum*, requires that we fight a war in such a manner that it leaves the door open for lasting peace after the conflict (Orend, 2002). That is, as history has shown, we should not brutalize an enemy, since that would leave ill-feelings to linger even after the fighting has stopped, which makes peaceful reconciliation most difficult to achieve. Robots do not necessarily represent an immoral or overly brutal way of waging war, but as they are needed for urban operations, such as patrolling dangerous streets to enforcing a curfew or securing an area, the local population may be less likely to trust and build good-will relationships with the occupying force (Sharkey, 2008). Winning hearts and minds is likely to require diplomacy and human relationships that machines would not be capable of delivering at the present time.

2.4.4. 'Comfort' Robots

Ethicists are already talking about the impact of robots as lovers or surrogate relationship partners (Levy, 2007). This does not seem so unthinkable, considering that some people already have 'relationships' with increasingly-realistic sex dolls, so robotics appear to be a natural next step in that industry; indeed, people today engage in sexual activities online, i. e., without a partner physically present.

In previous wars, women have been taken by the military to provide 'comfort' to soldiers or, in other words, forced into sexual slavery or prostitution. In World War II, women were most infamously used by the Japanese Imperial Army to satiate the pent-up carnal desires of its soldiers, ostensibly to prevent possible riots and discontent among the ranks; Nazi Germany reportedly also used women to stock their 'joy divisions' at labor or concentration camps. And instances of rape have been reported – and continue today – in armed conflicts from Africa to the Americas to Asia.

Robots, then, may be able to serve the same function of providing 'comfort' to the troops in a much more humane way, i. e., without the exploitation of women and prisoners of war. However, it is unclear that this function is truly needed (to the extent that most militaries today do not employ military prostitutes and seem to be operating adequately) or can overcome existing public inhibitions or attitudes on what is mostly a taboo subject of both sex in the military and sex with non-human objects.

2.5. Societal Challenges

2.5.1. Counter-Tactics in Asymmetric War

As discussed in the previous issue of lowering barriers to war or making war more risk-free, robots would help make military actions more effective and efficient, which is exactly the point of deploying those machines. Presumably, the more autonomous a robot is, the more lethal it can be (given requirements to discriminate among targets and so on). This translates to quicker, more decisive victories for us (your respective nation); but for the other side, this means swifter and perhaps more demoralizing defeats. We can reasonably expect that a consequence of increasing the asymmetry of warfare in our favor will cause opposing forces to engage in even more unconventional strategies and tactics, beyond ‘terrorist’ acts today (e. g., Kahn, 2002); few nations could hope to successfully wage war with militarily-superior states by using the same methods we use, given the sheer number of our troops and advanced military technologies and weapons.

This not only involves how wars and conflicts are fought, but also exposes our military as well as public to new forms of attack which may radically change our society. For instance, more desperate enemies may resort to more desperate measures, from intensifying efforts to acquire nuclear or biochemical weapons to devising a ‘scorched earth’ or ‘poison pill’ strategy that strikes deeply at us but at some great cost to their own forces or population (a Pyrrhic victory).

2.5.2. Proliferation

Related to the previous issue, history also shows that innovations in military technologies—from armor and crossbows to intercontinental missiles and ‘smart’ bombs—give the inventing side a temporary advantage that is eroded over time by other nations working to replicate the technologies. Granting that modern technologies are more difficult to reverse-engineer or replicate than previous ones, it nevertheless seems inevitable or at least possible that they can be duplicated, especially if an intact sample can be captured, such as immobilizing a ground robot as opposed to shooting down an unmanned aerial vehicle. So with the development of autonomous military robots, we can expect their proliferation with other nations at some future point. This means that these robots – which we are currently touting as lethal, difficult-to-neutralize machines – may be turned against our own forces eventually.

The proliferation of weapons, unfortunately, is an extremely difficult cycle to break: many nations are working to develop autonomous robotics, so a unilateral ban on their development would not accomplish much except to handicap that nation relative to the rest of the world. So the rush to develop this and other emerging technologies is understandable and irresistible, at least in today’s world. One possible defense for our pursuit, apart from self-interested reasons, is that we (again, your respective nation) want to ensure we develop these commanding technologies first, after which we would have more leverage to stop the proliferation of the same; further, because we occupy the higher moral ground, it would be most responsible for us to develop the technologies first.

The problem, of course, is that every nation thinks of itself as moral or ‘doing the right thing’, so it would be difficult to objectively assign a moral imperative to any given nation, including the US. Solving this problem, then, would seem to require

additional legal and ethical theorizing, likely resulting in new international treaties and amendments to the laws of war.

2.5.3. Space Race

As on earth, autonomous robots may hold many benefits for space exploration (Jónsson et al., 2007). Proliferation also has significant financial and environmental costs, particularly if military robotics technology is developed for outer space. First, launch costs are still astronomical, costing thousands of dollars per pound to put an object into low Earth orbit, and several times more per pound for geostationary orbit (not to mention periodic replacement costs and in-orbit repairs). An unlikely 'star wars' scenario aside – which would create countless pieces of space debris that would need to be tracked and threaten communications satellites and so on – even using robots for research purposes, e. g., to explore and develop moons and other planets, may spark another space race given the military advantages of securing the ultimate high ground. This not only opens up outer space for militarization, which the world's nations have largely resisted, but diverts limited resources that could make more valuable contributions elsewhere.

2.5.4. Technology Dependency

The possibility that we might become dependent or addicted to our technologies has been raised throughout the history of technology and even with respect to robotics. Today, ethicists worry that we may become so reliant on, for instance, robots for difficult surgery that humans will start losing that life-saving skill and knowledge; or that we become so reliant on robot for basic, arduous labor that our economy is somehow impacted and we forget some of those techniques (Veruggio, 2007). In the military, some soldiers already report being attached to the robot that saved their lives (Garreau, 2007).

As a general objection to technology, this concern does not seem to have much force, since the benefits of the technology in question often outweigh any losses. For instance, our ability to perform mathematical calculations may have suffered somewhat given the inventions of the calculator and spreadsheets, but we would rather keep those tools even at that expense. Certainly, it is a possible hypothetical or future scenario that, after relying on robots to perform all our critical surgeries, some event – say, a terrorist attack or massive electromagnetic pulse – could interrupt an area's power supply, disabling our machines and leaving no one to perform the surgery (because we forgot how and have not trained surgeons on those procedures, since robots were able to do it better). But as abilities enhanced by technology, such as performing numeric calculations, have not entirely disappeared from a population or even to a life-impacting degree in individuals, it is unclear why we would expect something as artful as brain or heart surgery to be largely lost. Similarly, in the case of relying on robots for manual labor, technology dependency would not erase our ability to, say, dig holes to plant trees to any impacting degree.

2.5.5. Civil Security and Privacy

Defense technologies often turn into public or consumer technologies, as we previously pointed out. So a natural step in the evolution of military robots would seem to be their incarnation as civil security robots; they might guard corporate buildings, control

crowds, and even chase down criminals. Many of the same concerns discussed above – such as technical challenges and questions of responsibility – would also become larger societal concerns: if a robot unintentionally (meaning that no human intentionally programmed it to ever do so) kills a small child, whether by accident (run over) or mistake (identification error), it will likely have greater repercussions than a robot that unintentionally kills a non-combatant in some faraway conflict. Therefore, there is increased urgency to address these military issues that may spill over into the public domain.

And while we take it that soldiers, as government property, have significantly decreased privacy expectations and rights, the same is not true of the public at large. If and when robots are used more in society, and the robots are likely to be networked, concerns about illegal monitoring and surveillance – privacy violations – may again surface, as they have with most other modern technologies, from DNA testing to genome sequencing to communications-monitoring software to nanotechnology. This raises the question of what kind of consent we need from the public before deploying these technologies in society.

2.6. Other and Future Challenges

2.6.1. Co-Opting of Ethics Effort by Military for Justification

A possible challenge that does not fit neatly into any of the above categories is the following political concern. Defense organizations may be aware (now) of the above concerns, but they may still not choose to address the issues to mitigate risk by absolving themselves of this responsibility: they may simply point to ethicists and robot scientists working on related issues as justification for proceeding ahead without any real plan to address at least some of these risks (Sharkey, 2007).

This is an interesting meta-issue for robot ethics, i. e., it is about the study and aims of robot ethics and not so much about an issue directly related to the use of autonomous robots. While it is certainly a possibility that organizations may only pay ‘lip-service’ to the project of robot ethics to appease critics and watchdogs, it does not take much enlightenment or foresight to see actual, real-world benefits from earnestly addressing these challenges. Further, we might measure the commitment that organizations have to robot ethics by the funding levels for such research. And it would be readily apparent if, for instance, defense organizations ignored the counsel and recommendations of experts engaged in the field. This is to say that co-opting is a relatively transparent activity to identify, although the point is more that it could be too late (for those harmed or society in general) by then.

2.6.2. Robot Rights

For now, robots are seen as merely a tool that humans use, morally no different (except in financial value) than a hammer or a rifle – their only value is instrumental, as a means to our ends. But as robots begin to assume aspects of human decision-making capabilities, the question may arise of their intrinsic value: do they deserve moral consideration of their own (beyond their financial or tactical value), and at what point in their evolution will they achieve this intrinsic value (as human lives seem to have)? When they become Kantian autonomous agents, making their own goals for themselves? Or would intrinsic value also require consciousness and emotions?

Some technologists have suggested that, by 2029, robots will demand equal treatment before the law with humans – and believe that this demand will be granted (e. g., Kurzweil, 1999). The only guarantee of avoiding this outcome appears to be a prohibition on programming robots with anything other than a ‘slave morality’, i. e., simply not allowing a Kantian-autonomous robot to ever be programmed or built (though such bans, especially when applied internationally, have been notoriously difficult to enforce). It will require careful consideration in the future as to whether such a prohibition should ever be lifted. Fortunately, even ‘technological optimists’, such as Kurzweil, do not expect this to be an issue until at least the 2020s.

Thus far, we have not discussed the possibility of giving rights to robots, not so much that it is farfetched to do so (e. g., we give rights to non-living entities such as corporations) or to consider them as persons (philosophically-speaking; e. g., again corporations or ships or some animals such as dolphins), but that the prerequisites for rights seem to require advanced software or artificial intelligence that is not quite within our foreseeable grasp. Specifically, if our notion of personhood specifies that only persons can be afforded rights and that persons must have free will or the capacity for free will, then it is unclear whether we will ever develop technologies capable of giving free will or full autonomy to machines, and, indeed, we don’t even know whether any other biological species will ever have or is now capable of such full autonomy; thus, we do not want to dwell on such a speculative issue here. That said, we will leave open the possibility that we may someday want or be logically required to give rights to robots (e. g., Kurzweil, 1999), but much more investigation is needed on the issue.

2.6.3. The Precautionary Principle

Given the above laundry list of concerns, some may advocate following a precautionary principle in robotics research – to slow or halt work until we have mitigated or addressed possible catastrophic risks – as critics have done for other technologies, such as bio- and nanotechnologies. For instance, those fearful of ‘Terminator’ scenarios where machines turn against us lesser humans, current research in autonomous robotics may represent a path towards possible, perhaps likely, disaster; thus a cautious, prudent approach would be to ban or at least significantly slow down research until we can sufficiently think about these issues before technology overtakes ethics. While we believe that a precautionary principle may be the appropriate course of action for some technology cases, many of the issues discussed above do not appear imminent enough to warrant a research moratorium or delay in the near term, just more investigation which may be sufficiently conducted in parallel to efforts to develop advanced robotics.

Furthermore, a cautionary approach in the development of advanced systems is inherently in tension with both the approaches taken by the scientists and engineers developing robots and with the outlook of military planners, rapidly searching for more effective tools for the task of waging war. We will again leave open the possibility that someday we may have to seriously consider the role of the precautionary principle in robotics, but that day appears to be in the distant horizon and does not demand an extensive discussion here.

3. Further Investigations Needed

Again, we do not intend the above to capture all possible issues related to risk and ethics in military robotics. Certainly, new issues will emerge depending on how the technology and intended uses develop. In the preceding, we have identified and started to discuss what we seemed to be the most urgent and important issues to resolve first, especially as related to responsibility, risk, and the ability of robots to discriminate among targets. This is only the beginning of a dialogue in robot ethics and merits further investigations.

For a more comprehensive discussion – that includes the current and predicted states of robotics, different programming approaches, the relevant laws of war, legal responsibility, product liability, and a framework for technology risk assessment – please see our complete report *Autonomous Military Robotics: Risk, Ethics, and Design* (Lin, Bekey, and Abney, 2008) at www.robotethics.com.

References

Anderson, M., & Anderson, S. L. (2007). Machine Ethics: Creating an Ethical Intelligent Agent. *AI Magazine* 28:4: 15–26.

Arkin, R. C. (2007). *Governing Lethal Behavior: Embedding Ethics in a Hybrid Deliberative/Hybrid Robot Architecture*. Report GIT-GVU-07-11, Atlanta, GA: Georgia Institute of Technology's GVU Center. Retrieved September 15, 2008 from: <http://www.cc.gatech.edu/ai/robot-lab/online-publications/formalization35.pdf>.

Asaro, P. (2007). *Robots and Responsibility from a Legal Perspective*. Proceedings of the IEEE 2007 International Conference on Robotics and Automation, Workshop on RoboEthics, April 14, 2007, Rome, Italy. Retrieved September 15, 2008, from: <http://www.peterasaro.org/writing/ASARO%20Legal%20Perspective.pdf>.

Asaro, P. (2008). How Just Could a Robot War Be? In Briggle, A., Waelbers, K., & Brey, P. (Eds.). *Current Issues in Computing and Philosophy* (pp. 50-64). Amsterdam, The Netherlands: IOS Press.

Asimov, I. (1950). *I, Robot* (2004 edition), New York, NY: Bantam Dell.

BBC (2005). *SLA Confirm Spy Plane Crash*. BBC.com. Retrieved September 15, 2008, from: http://www.bbc.co.uk/sinhala/news/story/2005/10/051019_uav_yavunia.shtml.

BBC (2007). *Robotic Age Poses Ethical Dilemma*. BBC.com. Retrieved September 15, 2008, from <http://news.bbc.co.uk/2/hi/technology/6425927.stm>.

Bekey, G. (2005). *Autonomous Robots: From Biological Inspiration to Implementation and Control*. Cambridge, MA: MIT Press.

Canning, J., Riggs, G. W., Holland, O. T., & Blakelock, C. (2004). *A Concept for the Operation of Armed Autonomous Systems on the Battlefield*, Proceedings of Association for Unmanned Vehicle Systems International's (AUVSI) Unmanned Systems North America, August 3-5, 2004, Anaheim, CA.

CBS (2007). *Robots Playing Larger Role in Iraq War*. Retrieved September 15, 2008, from <http://cbs3.com/topstories/robots.iraq.army.2.410518.html>.

Garreau, J. (2007). *Bots on the Ground*, Washington Post, May 6, 2007. Retrieved September 15, 2008, from: http://www.washingtonpost.com/wp-dyn/content/article/2007/05/05/AR2007050501009_pf.html.

Iraq Coalition Casualty Count (2008). *Deaths Caused by IEDs and U.S. Deaths by Month webpages*. Retrieved September 15, 2008, from: <http://icasualties.org/oif/IED.aspx> and <http://icasualties.org/oif/USDeathByMonth.aspx>.

Jónsson, A., Morris, R., & Pedersen, L. (2007). Autonomy in Space: Current Capabilities and Future Challenges, *AI Magazine* 28:4, 27-42.

Joy, B. (2000). Why the Future Doesn't Need Us, *Wired* 8:04, 238-262.

Kahn, P. (2002). The Paradox of Riskless War, *Philosophy and Public Policy Quarterly*, 22, 2-8.

Kurzweil, R. (1999). *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, New York, NY: Viking Penguin.

Kurzweil, R. (2005). *The Singularity is Near: When Humans Transcend Biology*, New York, NY: Viking Penguin.

Levy, D. (2007). *Love and Sex with Robots: The Evolution of Human-Robot Relationships*, New York, NY: HarperCollins Publishers.

Lin, P., Bekey, G., & Abney, K. (2008). *Autonomous Military Robotics: Risk, Ethics, and Design*, a report commissioned under US Department of the Navy, Office of Naval Research, award # N00014-07-1-1152, San Luis Obispo, CA: California Polytechnic State University. Retrieved May 10, 2009, from <http://www.roboethics.com>.

National Defense Authorization Act (2000). *Floyd D. Spence National Defense Authorization Act for Fiscal Year 2001*, Public Law 106-398, Section 220. Retrieved September 15, 2008, from: <http://www.dod.mil/dodgc/olc/docs/2001NDAA.pdf>.

National Transportation Safety Board (2007). NTSB Cites Wide Range of Safety Issues in First Investigation of Unmanned Aircraft Accident. NTSB press release, October 16, 2007. Retrieved September 15, 2008, from: <http://www.ntsb.gov/Pressrel/2007/071016b.htm>.

Orend, B. (2002). Justice After War. *Ethics & International Affairs*, 16:1, 43-56.

Padgett, T. (2008). *Florida's Blackout: A Warning Sign?* Time.com. February 27, 2008. Retrieved September 15, 2008, from: <http://www.time.com/time/nation/article/0,8599,1717878,00.html>.

Page, L. (2008). *US War Robots 'Turned Guns' on Fleshy Comrades*, The Register (UK), April 11, 2008. Retrieved September 15, 2008, from: http://www.theregister.co.uk/2008/04/11/us_war_robot_rebellion_iraq/.

Shachtman, N. (2007). *Robot Cannon Kills 9, Wounds 14*, Wired.com, October 18, 2007. Retrieved September 15, 2008, from: <http://blog.wired.com/defense/2007/10/robot-cannon-ki.html>.

Sharkey, N. (2007). *Automated Killers and the Computing Profession*, Computer 40, 122-124. Retrieved September 15, 2008, from: http://www.computer.org/portal/site/computer/menuitem.5d61c1d591162e4b0ef1bd108bcd45f3/index.jsp?&pName=computer_level1_article&TheCat=1015&path=computer/homepage/Nov07&file=profession.xml&xsl=article.xsl&.

Sharkey, N. (2008). *Cassandra or False Prophet of Doom: AI Robots and War*, IEEE Intelligent Systems, July/August 2008, pp. 14-17. Retrieved September 15, 2008, from: http://www.computer.org/portal/cms_docs_intelligent/intelligent/homepage/2008/X4-08/x4his.pdf.

Sofge, E. (2008). *The Inside Story of the SWORDS Armed Robot 'Pullout' in Iraq: Update*, PopularMechanics.com, April 15, 2008. Retrieved September 15, 2008, from: http://www.popularmechanics.com/blogs/technology_news/4258963.html.

Sparrow, R. (2007). Killer Robots, *Journal of Applied Philosophy*, Vol. 24:1, 62-77.

Thompson, P. B. (2007). *Food Biotechnology in Ethical Perspective*, 2nd ed., Dordrecht, The Netherlands: Springer.

US Army Surgeon General's Office (2006). *Mental Health Advisory Team (MHAT) IV: Operation Iraqi Freedom 05-07*, November 16, 2006. Retrieved September 15, 2008, from: <http://www.globalpolicy.org/security/issues/iraq/attack/consequences/2006/1117mhatreport.pdf>.

US Army Surgeon General's Office (2008). *Mental Health Advisory Team (MHAT) V: Operation Iraqi Freedom 06-08*, February 14, 2008. Retrieved September 15, 2008, from: http://www.armymedicine.army.mil/reports/mhat_v/Redacted1-MHATV-OIF-4-FEB-2008Report.pdf.

US Department of Defense (2007). *Unmanned Systems Roadmap 2007-2032*. Washington, DC: Government Printing Office. Retrieved September 15, 2008, from: <http://www.acq.osd.mil/USD/Unmanned%20Systems%20Roadmap.2007-2032.pdf>.

US Department of Energy (2004). *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations*. Washington, DC: Government Printing Office. Retrieved September 15, 2008, from: <https://reports.energy.gov/BlackoutFinal-Web.pdf>.

Van der Loos, H.F. M. (2007). *Ethics by Design: A Conceptual Approach to Personal and Service Robot Systems*, Proceedings of the IEEE Conference on Robotics and Automation, Workshop on Roboethics, April 14, 2007, Rome, Italy.

Veruggio, Giannarco (2007). EURON Roboethics Roadmap, Genova, Italy: European Robotics Research Network. Retrieved September 15, 2008, from: <http://www.roboethics.org/icra07/contributions/VERUGGIO%20Roboethics%20Roadmap%20Rel.1.2.pdf>.

Wallach, W. & Allen, C. (2008). *Moral Machines: Teaching Robots Right from Wrong*. New York, NY: Oxford University Press.

Walzer, M. (1977). *Just and Unjust Wars: A Moral Argument with Historical Illustrations*. New York, NY: Basic Books.

Weckert, J. (Ed.) (2007). *Computer Ethics*, Burlington, VT: Ashgate Publishing.